

# Section 1: Introduction

---

This section provides a brief overview of DASD performance evaluation, provides an overview of the DASD Component of CPEXpert, describes the sources of data used by the DASD Component to analyze system performance, and describes the performance data bases that the DASD Component can use.

## Chapter 1: Background

From a global view, DASD analysis can be viewed from two perspectives;

- How the DASD space is allocated (how much space is allocated versus how much space is free), how the allocated space is used (whether the allocated space is actually used by the files), and how frequently DASD files are referenced by applications (whether the files should be on-line or perhaps stored off-line). DASD performance analysis from this perspective typically relates to the cost of maintaining the on-line storage represented by the DASD devices.
- How the DASD configuration and files effect overall system performance or effect the performance of individual applications. DASD performance analysis from this perspective deals with the day-to-day operational performance implications of the DASD configuration and files.

There are numerous products which examine DASD performance from the first perspective mentioned above. CPEXpert does not depend upon these products, since a basic philosophy of CPEXpert is to use information normally available in a performance data base. Consequently, the DASD Component does not analyze DASD space allocation.

The DASD Component focuses on DASD performance from the second perspective: whether performance constraints exist in the basic DASD configuration and to what degree particular applications impact the performance of other, more important, applications.

Beginning with Release 12.2 of CPEXpert, the DASD Component also analyzes problems with VSAM data sets. This analysis is a partial automation of the analysis and guidance given in IBM's *VSAM Demystified* Redbook, SG24-6105.

Very sound DASD analysis advice was given by Friesenborg<sup>1</sup>: "It seems that the main thing to do is to avoid the outrageous; after that the subject doesn't warrant much attention." The major objective of DASD analysis is to identify and solve the "outrageous" or serious problems. The DASD Component of CPEXpert attempts to detect the serious problems, identify the causes of the problems, and suggest solutions to the problems.

---

<sup>1</sup>Friesenborg, S. E., IBM Publication GG22-9370

DASD performance constraints can be corrected by a variety of actions at different levels:

- At the file level (e.g., change file block size, move files to different location)
- At the device level (e.g. upgrade device)
- At the controller level (e.g., change the number of paths, upgrade controller, shorten string, add cache)
- At the system level (e.g., reconfigure paths)
- At the operations level (e.g., schedule conflicting applications to different times)
- At the applications level (e.g., change storage accessing requirements or patterns).

The DASD Component provides information and recommendations regarding which actions are suitable for specific problems in your environment.

## Chapter 2: The DASD Component of CPExpert

The DASD Component was implemented with the basic philosophy that:

- Most Data Base Administrators are able to address a limited number of DASD problems each day. The DASD Component was designed to identify the most serious DASD problems (both from the perspective of the overall DASD configuration and from the perspective of critical applications).
- Data Base Administrators are most interested in solving recurring DASD problems. The DASD Component was designed to identify the devices, files, and applications responsible for or experiencing recurring DASD problems.

With this design philosophy, the DASD Component (1) analyzes DASD information available in standard SMF/RMF records to select the most serious problems, and (2) identifies the devices and applications responsible for the problems or applications experiencing performance degradation because of the problems.

The guidance provided to the DASD Component allows a wide variety of options in the analysis process. These options allow users to define workload categories, to exclude specific DASD volumes from analysis, and to specify thresholds to limit the analysis.

These features and the analysis are provided by numerous SAS modules. These SAS modules (1) shape DASD performance and utilization data for detailed analysis by other modules, (2) determine the workload and analysis categories, (3) evaluate the data to assess potential areas of unacceptable performance, and (4) report the results from the evaluation. These modules are loaded and controlled by the central DASD Component of CPExpert (titled DASCPE).

- **Shape Data.** The majority of "processing" in the DASD Component is accomplished by numerous modules whose function is to combine, sort, summarize, correlate, and prepare the data for analysis. These modules extract and process relevant data elements from SMF records. Additionally, the modules incorporate a variety of queuing models to perform preliminary analysis of the data.
- **Determine workload and analysis categories.** A significant optional feature of the DASD Component is its ability to allow the specification of workload categories (e.g., TSO, CICS, Batch, etc.) and to focus analysis with respect to these workload categories. Additional controls provide the flexibility to exclude DASD volumes from consideration.
- **Evaluate Data.** The evaluation of the data is accomplished by rules whose purpose is to determine whether DASD performance was acceptable, and to analyze the DASD devices with unacceptable performance. These rules determine the devices with major problems, identify the probable causes of the problems, and (if expanded analysis is performed) identify the applications most likely responsible for the problems.

- **Report Results.** The DASD Component reports the results from the evaluation, in a variety of reports at the device or application workload level.

Each significant finding is described in Appendix A of this document. The description summarizes the finding, lists predecessor findings, discusses the rationale for the finding, and recommends action.

- The summary presents a short description of the finding.
- The predecessor findings are listed so you can follow the line of reasoning leading to a particular rule being executed.
- The discussion describes as much as necessary of the operation of the computer system as it relates to the particular finding. The purpose of the discussion is to explain the reasoning behind the finding. If appropriate, the discussion might refer you to related discussions in the DASD Component User Manual, or in a User Manual of another CPExpert component (e.g., the User Manual for the WLM Component may be used as a reference to avoid repeating a detailed discussion).
- The recommendations suggest possible actions that should be considered based on the findings. In many cases, multiple possible actions are listed. You must determine which actions should be taken (this determination is based upon the suitability of the actions to your own environment, the financial implications of the action, and the "political" acceptability of the action.)

## Chapter 3: Data Sources

CPEXpert analyzes the performance of the DASD configuration based upon data from three sources:

- **RMF information.** The basic analysis performed by the DASD Component relies on standard SMF Type 70(series) information. The DASD Component analyzes Type 70, Type 72, Type 74, Type 75, and Type 78CF information to build a model of the I/O configuration in the sysplex, and to analyze performance of devices that performed poorly relative to other devices in the configuration.
- **SMF information.** The DASD Component of CPEXpert optionally includes a modification to MXG or to NeuMICS (using standard exit facilities). The modification allows extraction of SMF Type 30(Data Definition) information related to the individual DASD devices used by job steps . This information allows the DASD Component to perform "expanded" analysis, analyzing DASD performance based upon critical workloads.

Additionally, the DASD Component optionally extracts data from SMF Type 42 (Data Set Statistics) to report information about the data sets residing on devices with poor performance.

- **Guidance information.** Guidance information is contained in CPEXPERT.USOURCE(DASGUIDE). This PDS member contains information that allows you to specify specific time periods in which CPEXpert should analyze data, which performance groups (for Compatibility Mode) or service classes (for Goal Mode) are associated with different workload categories, what level of detail should be pursued in the DASD analysis, which volumes should be excluded from analysis, device path and caching characteristics, and other variables to control the operation of the DASD Component.

## Chapter 4: Performance Data Bases

The "raw" RMF data contained in the SMF Type 30(DD) and Type 70(series) records must be translated into SAS format and placed into a performance data base before CPEXpert can use the information. There are two ways in which these records can be placed into a performance data base:

- **MXG performance data base.** The performance data base can be created by Merrill's Expanded Guide (MXG) software. MXG is provided by Merrill Consultants, Dallas, Texas. MXG provides a low-cost mechanism by which installations can create and maintain a performance data base.
- **NeuMICS performance data base.** The performance data base can be created by the MVS Integrated Control System (NeuMICS). NeuMICS is provided by Computer Associates, Vienna, Virginia. The NeuMICS SMF (or Base) Component creates and maintains a comprehensive performance data base from SMF/RMF information.

The flexibility to use either of the above performance data bases is due to the fact that CPEXpert is implemented in the SAS language. SAS provides a powerful macro facility, both with respect to macro coding and with respect to macro variable names.

CPEXpert generally uses SAS macro variable names when referring to an element of information in the SMF records. CPEXpert uses the SAS "%LET" statements to define the macro variables as MXG variable names or NeuMICS variable. These "%LET" statements are contained in unique variable definition modules for MXG and NeuMICS.

Thus, the same CPEXpert software can be executed against either of the performance data bases, by only invoking the proper definition module. The SAS %LET statements in the definition module automatically cause CPEXpert to refer to the proper MXG or NeuMICS data elements.

## Chapter 5: Types of Analysis

The DASD Component performs two types of analysis of DASD performance and constraints to improved performance: **basic** analysis and **expanded** analysis. Additionally, the expanded analysis can analyze DASD performance from the perspective of specific application workloads or from the perspective of specific data sets. With CPExpert Release 12.2, the DASD Component analyzes potential performance problems with VSAM data sets. This section describes these types of analysis.

### Chapter 5.1: Basic Analysis

The philosophy of the basic analysis is to identify the devices which have the most potential for improvement, and to analyze these devices in detail. The method used to identify a candidate set of devices is quite simple:

- CPExpert computes the average device response time for each **type** of device in the configuration, for each RMF measurement interval. The logic computes the average device response by type of device, since better performance would be expected from cached devices (for example) than from non-cached devices. This method essentially assesses the performance of each device against the performance of similar devices in the configuration.
- Devices that exceed the average device response time for their device type in any RMF measurement interval are selected as candidates for improvement. The rationale is that improvement efforts should not be directed at devices which provide better than average response. Thus, the candidate set of devices to analyze consists of those which provided worse than average response.
- The I/O rate of each candidate device is weighted by its response time. The result is a measure of the relative performance improvement **potential** of each device which provided worse than average response, from an overall system view. For example, consider two devices in a device type having an average I/O response of 20 milliseconds:

Device A: I/O rate = 30 I/O operations per second  
Device response = 25 milliseconds  
Weighting factor =  $30 * 25 = 750$

Device B: I/O rate = 5 I/O operations per second  
Device response = 40 milliseconds  
Weighting factor =  $5 * 40 = 200$

In the above example, CPExpert would select Device A as having the most overall potential for improvement, even though its per-I/O device response was not as bad as the device response of Device B.

This logic totally ignores the potential situation where Device B holds critical data sets (an index, for example) and thus has more effect on overall system performance than the I/O rate indicates. Unfortunately, CPExpert is examining measurement data and has no way to assess relative importance of devices. (In the basic analysis, CPExpert **does** provide an option by which you can select specific devices, and CPExpert will analyze performance constraints of only those devices.)

- A model is built of the I/O configuration (using Type 70, Type 73, Type 74, Type 75, and Type 78CF information). Measurement data (channel, controller, device information) is processed against this model. CPExpert uses queuing models to compute the likely delays at various physical or logical parts of the configuration (for example, path delays, missed RPS reconnect delays, etc.). The result from this analysis associates probable delays with the specific devices.
- The device(s) with the overall most potential for improvement is selected for detailed analysis as to likely causes (seeking, missed RPS reconnect, etc.).
- The results from the above analysis are reported, normally in a one or two page report.

## Chapter 5.2: Expanded Analysis (Specific applications)

CPEXpert can perform an expanded analysis of DASD performance from the perspective of specific applications. This section describes how this analysis is accomplished.

- A modification is made to the MXG or NeuMICS software creating your performance data base. This modification normally uses standard user exits<sup>2</sup>, adds an insignificant amount of processing time, and uses a small amount of DASD space.
- The modification normally is invoked as MXG or NeuMICS process the SMF Type 30(Interval) records.

For MXG, the modification **could** be invoked as MXG processes the SMF Type 30(Job Step Termination) records if you do not collect SMF Type 30(Interval) records. Using Job Step Termination records is discouraged since incomplete analysis may result (the job steps may span many RMF measurement intervals and it may be impossible to correlate the data recorded by RMF with lengthy job steps).

- The modification records the system, job, job step, performance group or service class, and summary information about each job step's use of DASD devices by device number. The information is recorded only at the device level (rather than at the DD name level) and only required information is retained. Consequently, the records are small and require a small amount of DASD space.
- You define specific applications (or "loved one" applications) to CPEXpert. Specific workload categories are defined (for example, TSO or CICS), and you define the performance groups or service classes assigned to the workload categories.
- CPEXpert processes the data acquired using the modification to MXG or NeuMICS described above, and matches application use of devices against SMF Type 74 device information. The resulting information is application-specific, in that it relates device use to the applications using the devices.
- CPEXpert then identifies the devices used by "loved one" workloads, computes the average device response from this set of devices, and selects all devices exceeding the average response as candidates for analysis.
- The normal basic analysis is performed of all devices resulting from the above selection criteria. The result is an analysis from the perspective of specific applications.

---

<sup>2</sup>A single one-line modification is required to NeuMICS code in the NeuMICS DYSMFFMT module. NeuMICS does not currently provide an exit which can be used to add a data set to the output from NeuMICS as it analyzes SMF data. Such an exit does exist for the NeuMICS CICS Component (the \_USRSDKP exit), and a request has been made that a similar exit be provided for the NeuMICS BASE Component.

- Additionally, CPExpert can identify non-critical applications causing performance problems for the "loved one" workload. For example, CPExpert can identify all non-critical applications using a significant amount of a path if CPExpert discovers that path contention results in poor performance to "loved one" applications.

## Chapter 5.3: Expanded Analysis (Specific data sets)

CPEXpert can perform an expanded analysis of DASD performance from the perspective of specific data sets. The analysis is done in one of two ways: (1) analyzing data set response based on information from TYPE42DS statistics or (2) analyzing data based on information in TYPE14/15 and using the CPEXpert modification to MXG or NeuMICS as the SMF Type30 records are processed.

- **Analysis based on TYPE42DS.** This method can provide comprehensive information, without requiring the modification to MXG. This method **is not applicable** to NeuMICS, since NeuMICS does not retain sufficient information related to SMF Type 42 (Data Set Statistics).

MXG creates TYPE42DS from the SMF Type 42 (Subtype 6 - data set I/O statistics) records created by SMS. The TYPE42DS file contains I/O access characteristics information, at the data set level. CPEXpert extracts data set information and data set response statistics from TYPE42DS, and compares these to the response objectives you have made in USOURCE(DASGUIDE) about critical data sets. When any data set response time exceeds the specified objective, the DASD Component selects the data set and the volume it resides on for detailed analysis.

- **Analysis based on TYPE14/15 and CPEXpert modification.** This method is a bit more involved, but can be used regardless of whether data sets are managed by SMS, and this method can be used regardless of whether you use MXG or NeuMICS to create your performance data base.

With this method, the DASD Component analyzes SMF Type 14/15 records to extract data set names that correspond to the critical data sets that you have identified in USOURCE(DASGUIDE). The CPEXpert modification to MXG or NeuMICS extracts DD information as SMF Type 30 records are processed. The DASD Component then correlates the data set information with the DD information, to determine whether critical data sets exceed the specified response objective. When any data set responses time exceeds the specified objective, the DASD Component selects the data set and the volume it resides on for detailed analysis.

With either of the above methods, CPEXpert matches the data set information against SMF Type 74 device information, and identifies each device which has an average I/O response time exceeding the response objective of any data set residing on the device. The resulting information is specific to devices which have I/O response times exceeding a response objective.

The I/O rate of each device identified above is weighted by its response time. The result is a measure of the relative performance improvement potential of each device which provided I/O response exceeding the response objective for some data set on the device.

The normal basic analysis is performed of all devices resulting from the above selection criteria. The result is an analysis from the perspective of specific data sets.

Additionally, CPExpert creates a report describing the I/O activity of all data sets which resided on devices having an I/O response worse than the response objective for the data set.

## Chapter 5.4: Analysis of shared DASD conflicts

CPEXpert can perform an analysis of conflicts between DASD shared by systems or MVS images. This section describes how this analysis is accomplished.

Assume that CPEXpert is analyzing DASD performance problems with System A, and some DASD devices are shared with System B. There are three major ways in which the shared DASD can cause performance problems for System A:

- **Case #1.** The arm of the device has been moved by System B when System A attempts to access a cylinder. The DASD I/O operation from System A must move the arm to the desired cylinder, as a SEEK operation.

CPEXpert computes the time required to perform SEEKS on System A. If the computed SEEK time is a major cause of performance problems, CPEXpert analyzes the data from System B to determine whether System B generates a large number of I/O operations to the device.

- If System B does **not** generate a relatively large number of I/O operations to the device, CPEXpert concludes that there is **not** a conflict. There is little doubt about the validity of this conclusion. If System B does not direct much I/O activity to the device, System B clearly cannot cause seek problems for System A.
- If System B **does** generate a relatively large number of I/O operations to the device, CPEXpert concludes that there **is** a conflict caused by sharing the device.

To be absolutely correct, CPEXpert should process the configuration definitions for System B, process System B's channel and device information, and compute seek information for System B. CPEXpert could then determine whether System B also experienced a high seek rate for the device.

If both System A and System B experienced a high seek rate, CPEXpert could be absolutely sure that there was a shared DASD conflict. This approach would unnecessarily use system resources and would be cumbersome to implement.

Consequently, CPEXpert makes the assumption that I/O operations to the device are random between System A and System B. CPEXpert thus can conclude that if System A experiences a high seek rate and System B significantly uses the device (exhibited by a high I/O rate), then System B must also experience a high seek rate. To assume otherwise would require that I/O from System B be coordinated with the I/O from System A, such that System B does not experience seeking similar to System A.

Since System A experiences a high seek rate and System B experiences a high I/O rate, CPEXpert concludes that there is a conflict caused by the shared DASD.

- **Case #2.** The device (or the path to the device) is busy to System B when System A attempts to reference the device. The DASD I/O operation from System A must wait until the device (or path) is free. This wait time is directly reflected as PEND time in the RMF measurement data.

CPEXpert determines whether PEND time is a major cause of performance problems for the device (or devices) selected for detailed analysis. If PEND time is a major cause of performance problems, CPEXpert analyzes the data from System B to determine whether DASD I/O operations from System B could cause the PEND delays for System A.

CPEXpert analyzes the total CONN time and DISC time for the device from other systems sharing the device. CPEXpert concludes that the PEND delay is caused by other systems if the total CONN time and DISC time is more than 25% of the PEND delay for the device experienced on System A.

- **Case #3.** The device is cached, the cached controller is shared between System A and System B, and I/O operations from System B preempt data in the cache which is required by System A.

This problem is evidenced by relatively large I/O times for System A. Since the data required by System A is not in the cache, the data must be acquired from the device. This potentially results in a seek operation to position the arm, and will require normal latency and data transfer from the device.

In any of the above situations with cached shared devices, CPEXpert analyzes the amount of I/O operations from System A and System B to the cached controller.

- If System B does **not** generate a relatively large number of I/O operations to the device, CPEXpert concludes that there is **not** a conflict.
- If System B **does** generate a relatively large number of I/O operations to the device, CPEXpert concludes that there **is** a conflict caused by sharing the device. The rationale for these conclusions is the same as was discussed under Case #2 (seek conflicts).

**The standard analysis performed by CPEXpert may not detect a cache problem under two possible scenarios.**

- If only one cached controller is attached to System A, CPEXpert may not detect a problem with the device. This is because the logic employed by CPEXpert selects devices with the most performance improvement within each **type of device** and then selects the **overall** "worst" devices for detailed analysis.

CPEXpert considers cached devices to be a unique type. If all devices on the cached controller received bad service caused by shared cached problems,

CPEXpert may not detect a performance problem with the cached devices. This is because all devices in the "device type" could have roughly equal poor service and no device would be **significantly** worse than the other devices in the device type. Consequently, CPEXpert might not select any of the cached devices for detailed analysis

If there are multiple cached controllers, there would be a larger number of "candidate" devices, and the standard analysis performed by CPEXpert is more likely to identify any problem caused by shared cache controllers.

- The analysis performed by CPEXpert may not detect a cache problem if the cache is being replaced with data from another volume. It is possible that System B could cause data from another volume to be loaded into the controller's cache. This volume could be a volume not be flagged as the "worst" performing volume when analyzing performance from the perspective of System A. From System A's perspective, accesses to the worst volume would simply not find required data in cache.

CPEXpert cannot at present relate (1) poor performance for one volume on System A caused by cache problems and (2) I/O operations to a different volume by System B.

The shared DASD analysis is accomplished when you perform the following steps:

- You collect and process SMF data for each system to be analyzed. The SMF data must be placed into your standard performance data base. This step is not specific for CPEXpert; it is the normal processing you would do. The step is listed first just for completeness.
- **You should ensure that the system clocks on all systems are set at roughly the same time.** CPEXpert will examine the SMF data contained in your performance data base, based upon the SMF time associated with the measurement data. Shared DASD problems between systems will be detected based on this SMF time. If the SMF times are significantly different, the analysis might not properly identify conflicts or conflicting applications.

You should not worry that the system clocks be set at **exactly** the same time. The analysis performed by CPEXpert is not intended to identify an isolated performance problem. Rather, CPEXpert attempts to identify those problems which **continually** cause shared DASD performance problems. If a shared DASD problem is continual, the problem will be reflected in multiple SMF recording intervals. Consequently, it is not essential that the SMF times be exactly matched between systems.

- You specify **%LET SHARED = YES;** in USOURCE(DASGUIDE) to tell CPEXpert to perform analysis of shared DASD.

CPEXpert performs the following processing if you have indicated that an analysis of potential conflicts between shared DASD should be performed:

- CPExpert determines whether the "worst" device selected for detailed analysis is attached to a control unit shared with another system. If so, CPExpert performs an analysis of potential conflicts caused by shared DASD.
- CPExpert identifies other systems which reference the "worst" device by analyzing the SMF Type 74 data in the performance data base relating to all other systems. The SMF Type 74 data contain the VOLSER for each device referenced. CPExpert simply selects SMF Type 74 information for the systems which reference the VOLSER of the "worst" device. This information is retained for more detailed analysis about potential conflicts.
- Once CPExpert has identified all systems which reference the "worst" device, CPExpert analyzes the DASD I/O characteristics of these systems with respect to the "worst" device. As described earlier, the analysis makes a basic assumption that the I/O activity from the different systems is random among the systems (that is, the I/O activity of System B is independent from the I/O activity of System A).

Consequently, CPExpert can conclude that delays on System A are caused by shared DASD conflicts only if (1) System B experienced the same delays with the "worst" device (in the case of PEND delays) or (2) System B generated a high I/O activity to the "worst" device. CPExpert can conclude that delays on System A are not caused by shared DASD conflicts if the data from System B does not exhibit these characteristics. The earlier discussion explains the rationale for these conclusions.

- If "expanded" analysis is being performed, CPExpert can identify the applications on System B which reference the "worst" device.

SMF Type 30 records do not contain the VOLSER in the DD section. However, CPExpert correlates the information acquired from other records to relate the VOLSER to the device number. While extracting data from the SMF Type 74 information in the performance data base, CPExpert retains the device number associated with the devices. This device number is matched against the device numbers in the SMF Type 30 records to identify the VOLSER associated with the device.

CPExpert thus can identify all applications from System B which reference the "worst" device shared with System A.

## Chapter 5.5: Analysis of VSAM data sets

CPEXpert can perform an analysis of problems with VSAM data sets, **if MXG has created the performance data base**. This section describes how this analysis is accomplished.

VSAM data set activity typically accounts for a large percent of I/O activity (more than 70% at some sites). Tuning of a few files or correcting common problems often can result in significantly improved performance (IBM benchmarks show up to 90% improvement resulting from some simple changes).

With CPEXpert Release 12.2, the DASD Component has been enhanced to provide a rudimentary analysis of common VSAM problems. Additional analysis will be added in future enhancements to the DASD Component.

The DASD Component can optionally analyze VSAM data set performance problems or potential problems under the following conditions:

- MXG has created the performance data base. Unfortunately, there is insufficient data retained in a NeuMICS performance data base to allow CPEXpert to analyze VSAM performance problems.
- The performance data base contains the MXG TYPE64 file (and CPEXpert is provided with the **%LET TYPE64=Y**; guidance variable).
- Most analysis of VSAM data sets also requires the MXG TYPE64 file (and CPEXpert is provided with the **%LET TYPE64=Y**; guidance variable). This is because the analysis depends on being able to identify the VSAM file type (e.g., KSDS, VRRDS, etc.), and the buffering technique used (e.g., NSR or LSR).